

High-Dimensional Multiple Testing and Multi-Armed Bandits

Robert Nowak

University of Wisconsin-Madison

Abstract

This lecture will discuss relations between large-scale multiple testing and multi-armed bandit (MAB) problems. The limits of detection in the “sparse normal means” testing problem are well understood [c.f. Donoho and Jin, '04]. In this talk, I will show that sequential testing procedures (a simple form of adaptive/active sampling) can provide significant improvements. Moreover, in many applications, the sparse means model is too restrictive. I will discuss the classic and more general “best population” problem [Bechhofer, '54], which considers sampling from n different distributions with means μ_1, \dots, μ_n in $[0, 1]$ in order to find the distribution(s) with the largest mean(s). This is called the “best arm” problem in the MAB literature. I will discuss adaptive sampling strategies for the problem and bounds on the total number of samples needed to identify the distribution with the largest mean. The bounds depend on the underlying means and thus the sampling strategies automatically adapt to specific instance of the problem. In addition to the theory, the talk will also present some real-world applications.